arXiv:1603.05609v2 [cs.CG] 15 Nov 2016

ABCDepth: efficient algorithm for Tukey depth

Milica Bogićević* and Milan Merkle†

University of Belgrade, Faculty of Electrical Engineering, Bulevar Kralja Aleksandra 73, 11120 Belgrade, Serbia

Abstract. We present a new fast approximate algorithm for Tukey (halfspace) depth level sets and its implementation. Given a d-dimensional data set for any $d \geq 2$, the algorithm is based on a representation of level sets as intersections of balls in $\mathbf{R}^d$ (M. Merkle, J. Math. Anal. Appl. **370** (2010)). Our approach does not need calculations of projections of sample points to directions. This novel idea enables calculations of level sets in very high dimensions with complexity which is linear in d, which provides a great advantage over all other approximate algorithms. Using different versions of this algorithm we demonstrate approximate calculations of the deepest set of points ("Tukey median"), Tukey's depth of a sample point and of out-of-sample point as well as approximate level sets that can be used for constructing depth contours, all with a linear in d complexity. An additional theoretical advantage of this approach is that the data points are not assumed to be in "general position". Examples with real and synthetic data show that the executing time of the algorithm in all mentioned versions in high dimensions is much smaller than other implemented algorithms and that it can accept thousands of multidimensional observations.

Keywords: Big data, multivariate medians, depth functions, computing Tukey's depth.

1. Introduction

A basic statistical task is to simplify a large amount of data using some values derived from the data set as representative points. Among many ways to choose representative points, a natural idea is to choose those that are located in the center of the data set. One way to define a center is to define what is meant by deepness, and then to define the center as the set of deepest points.

Although this paper is about multivariate medians and related notions, for completeness and understanding some ideas, we start from the univariate case. Talking in terms of probability distributions, let X be a random variable and let $\mu = \mu_X$ be the corresponding distribution, i.e., a probability measure on $(\mathbf{R}, \mathcal{B})$ so that $P(X \leq x) = \mu\{(-\infty, x]\}$. For univariate case, a median of X (or a median of μ_X) is any number m such that $P(X \leq m) \geq 1/2$ and $P(X \geq m) \geq 1/2$. In terms of data set, this property means that to reach any median point from outside of the data set, we have to pass at least $1/2$ of data points, so this is the deepest point within the data set. With respect to this definition, we can define the depth of any point $x \in \mathbf{R}$ as

$$D(x, \mu) = \min\{P(X \leq x), P(X \geq x)\} = \min\{\mu((-\infty, x]), \mu([x, +\infty))\}. \tag{1}$$

The set of all median points $\{\text{Med}\mu\}$ is a non-empty compact interval (can be a singleton). It can be shown that (see [15], [16])

$$\{\text{Med}\mu\} = \bigcap_{\text{J=[a,b]: } \mu(\text{J})>1/2} \text{J}, \tag{2}$$

and (2) can be taken for an alternative (equivalent) definition of univariate median set. In $\mathbf{R}^d$ with $d > 1$, there are quite a few different concepts of depth and medians (see for example [23], [26], [30]). In this paper we propose an algorithm for halfspace depth (Tukey's depth, [28]), which is based on extension and generalization of (2) to $\mathbf{R}^d$ with balls in place of intervals as in [16].

The rest of the paper is organized as follows. Section 2 deals with a theoretical background of the algorithm in a broad sense. In Section 3 we present approximate algorithm for finding Tukey median as well as versions of the same algorithm for finding Tukey depth of a sample point, the depth of out-of-sample point, and for data contours. We also provide a derivation of complexity for each version of the

*antomripmuk@yahoo.com

†emerkle@etf.rs

algorithm and present examples. Section 4 provides a comparison with several other algorithms in terms of performances.

2. Theoretical background: Depth functions based on families of convex sets

Definition 2.1. *Let $\mathcal{V}$ be a family of convex sets in $\mathbf{R}^d$, $d \geq 1$, such that: (i) $\mathcal{V}$ is closed under translations and (ii) for every ball $B \in \mathbf{R}^d$ there exists a set $V \in \mathcal{V}$ such that $B \in \mathcal{V}$. Let $\mathcal{U}$ be the collection of complements of sets in $\mathcal{V}$. For a given probability measure μ on $\mathbf{R}^d$, let us define*

$$D_{\mathcal{V}}(x;\mu) = \inf\{\mu(U) \mid x \in U \in \mathcal{U}\} = 1 - \sup\{\mu(V) \mid V \in \mathcal{V},\ x \in V'\} \tag{3}$$

The function $x \mapsto D(x;\mu,\mathcal{V})$ will be called a depth function based on the family $\mathcal{V}$.

Remark 2.1. Definition 2.1 is a special case of Type D depth functions as defined in [30] which can be obtained by generalizations of (1) to higher dimensions. The conditions stated in [16] that provide desirable behavior of the depth function, are satisfied in this special case, with additional requirements that sets in $\mathcal{V}$ are closed or compact.

Example 2.1. 1° Let $\mathcal{V}$ be the family of all compact intervals $[a,b] \subset \mathbf{R}$. As shown in [16], the depth function based on $\mathcal{V}$ is the same as the one defined by (1).

2° With $d = 2$, consider the family $\mathcal{V}$ of rectangles with sides parallel to coordinate axes. The corresponding depth function reaches its maximum $D_{\max} \geq 1/2$ at coordinate-wise median. The same holds for $d > 2$, with "boxes" whose sides are parallel to coordinate hyper-planes.

3° For $d > 1$, let K be a closed convex cone in $\mathbf{R}^d$, with vertex at origin, and suppose that there exists a closed hyperplane π, such that $\pi \cap K = \{0\}$ (that is, $K \setminus \{0\}$ is a subset of one of open halfspaces determined by π). Define a relation $\preceq$ by $\boldsymbol{x} \preceq \boldsymbol{y} \iff \boldsymbol{y} - \boldsymbol{x} \in K$. Generalized intervals based on this partial order can be defined as

$$[\boldsymbol{a},\boldsymbol{b}] = \{\boldsymbol{x} \mid \boldsymbol{x} - \boldsymbol{a} \in K \wedge \boldsymbol{b} - \boldsymbol{x} \in K\} = (\boldsymbol{a} + K) \cap (\boldsymbol{b} - K).$$

Now let us take $\mathcal{V}$ to be a collection of all such intervals with finite endpoints and define the depth by (3). It can be shown ([16, Section 3]) that the maximal depth is always $\geq 1/2$, and the median set can be found using formula (2) with generalized intervals.

4° Let us consider a family $\mathcal{V}$ of all closed halfspaces in $\mathbf{R}^2$. Let X be a random point in $\mathbf{R}^2$ with $P(X = A) = P(X = B) = P(X = C) = 1/3$, where ABC is a non-degenerated triangle. Here all points inside and on the border of the triangle ABC have the depth $1/3$ and the depth of other points is equal to zero. Similar examples can be made for arbitrary dimension (see Example 4.1. in [16]).

From the above examples we see that

- A family $\mathcal{V}$ is not uniquely determined by the depth function; if we start with different collections $\mathcal{V}_1$ and $\mathcal{V}_2$, the corresponding depth functions based on them can be the same (see [16, Theorem 4.2] for a set of sufficient conditions).
- The maximal depth in $d > 1$ doesn't need to be $1/2$ as in the scalar case. The following general result ([16, Theorem 4.1], see also particular case for Tukey's depth in [22, Proposition 9]) says that the maximal depth has to be $1/(d+1)$ or bigger:

Theorem 2.1. *Let $\mathcal{V}$ be any non-empty family of compact convex subsets of $\mathbf{R}^d$ satisfying the conditions as in Definition 2.1. Then for any probability measure μ on $\mathbf{R}^d$ there exists a point $x \in \mathbf{R}^d$ such that $D_{\mathcal{V}}(x;\mu) \geq \frac{1}{d+1}$.*

The set of points with maximal depth is called *the center of distribution* and denoted as $C(\mu,\mathcal{V})$. In general, one can observe *level sets* (or *depth regions* or *depth-trimmed regions*) of level α which are defined by

$$S_\alpha = S_\alpha(\mu,\mathcal{V}) := \{x \in \mathbf{R}^d \mid D(x;\mu,\mathcal{V}) \geq \alpha\}. \tag{4}$$

Clearly, if $\alpha_1 < \alpha_2$ then $S_{\alpha_1} \supseteq S_{\alpha_2}$ and $S_\alpha = \emptyset$ for $\alpha > \alpha_m$, where α_m is the maximal depth for given probability measure μ.

The borders of depth level sets are called *depth contours* (in two dimensions) or *depth surfaces* in general. Let us note that the all statistical inference based on multidimensional depths is performed using level sets and contours (see [7, 8, 29, 22]), and that it is rarely necessary to find a depth of a particular point. On the other hand, in order to describe level sets and the center of distribution we do not need to calculate depth functions, as the next result shows ([30] and Theorem 2.2. in [16]).

Theorem 2.2. *Let $D(x;\mu,\mathcal{V})$ be defined for $x \in \mathbf{R}^d$ as in Definition 2.1. Then for any $\alpha \in (0,1]$*

$$S_\alpha(\mu,\mathcal{V}) = \bigcap_{V\in\mathcal{V},\mu(V)>1-\alpha} V. \tag{5}$$

The center of a distribution is then the smallest non-empty level set; equivalently,

$$C(\mu,\mathcal{V}) = \bigcap_{\alpha:S_\alpha\neq\emptyset} S_\alpha(\mu,\mathcal{V}) \tag{6}$$

Since sets in $\mathcal{V}$ are convex, the level sets are also convex.

From (4) and (5) we can see that the depth function can be uniquely reconstructed starting from level sets.

Corollary 2.1. *For given μ and $\mathcal{V}$, let S_α, $\alpha \geq 0$ be defined as in (5), with $S_\alpha = \emptyset$ for $\alpha > 1$. Then the function $D : \mathbf{R}^d \mapsto [0,1]$ defined by*

$$D(x) = h \quad \Longleftrightarrow \quad x \in S_\alpha, \alpha \leq h \quad \text{and} \quad x \notin S_\alpha, \alpha > h \tag{7}$$

is the unique depth function such that (4) holds.

The algorithm that we propose in this paper primary finds level sets based on the formula (5), rather then directly the depth of particular points. The depth of a single point, if needed, can be calculated via Corollary 2.1. The algorithm will be demonstrated in the case of half-space depth, which is described in the next section.

3. ABCDepth Algorithm for Tukey depth: Implementation and the Output

The most popular choice among depth functions of Definition 2.1 is the one which is based on half-spaces, also called Tukey's depth [28]. Here $\mathcal{V}$ is the family of all open half-spaces, and the complements are closed half-spaces, so the usual definition of Tukey depth is obtained from (3) as

$$D(x;\mu) = \inf\{\mu(H) \mid x \in H \in \mathcal{H}\}, \tag{8}$$

where $\mathcal{H}$ is the family of all closed halfspaces. In this section, we consider only half-space depth, so we use the notation $D(x,\mu)$ instead of $D_{\mathcal{V}}(x,\mu)$. As already noticed in Section 2, a depth function can be defined based on different families $\mathcal{V}$. We say that families $\mathcal{V}_1$ and $\mathcal{V}_2$ are depth-equivalent if $D_{\mathcal{V}_1}(x;\mu) = D_{\mathcal{V}_2}(x;\mu)$ for all $x \in \mathbf{R}^d$ and all probability measures μ. Sufficient conditions for depth-equivalence are given in [16, Theorem 2.1], and it was shown there that in the case of half-space depth the following families are depth-equivalent: a) Family of all open halfspaces; b) all closed halfspaces; c) all convex sets; d) all compact convex sets; e) all closed or open balls.

For determining level sets we choose closed balls, and so we can define $\mathcal{V}$ as a set of all closed balls (hyper-spheres) and (exact) level sets can be found as

$$S_\alpha(\mu,\mathcal{V}) = \bigcap_{B\in\mathcal{V},\mu(B)>1-\alpha} B. \tag{9}$$

From now on we consider only the case when the underlying probability measure μ is derived from a given data set.

3.1. **The sample version.** In the setup with data sets, we have a sample of n points $\{x_1, \dots, x_n\}$ (with repetitions allowed) and we may use the counting measure defined as

$$\mu(A) = \frac{\#\{x_i : \ x_i \in A\}}{n}. \tag{10}$$

In this part we assume to have a fixed sample of n points, so we don't need to explicitly acknowledge the dependence of sample and its cardinality. The level sets in (9) for $\alpha \in (0, 1]$ can be found as

$$S_\alpha = \bigcap_{B \in \mathcal{V}, \#\{x_i : \ x_i \in B\} \geq \lfloor n(1-\alpha)+1 \rfloor} B, \tag{11}$$

where we write S_α instead of $S_\alpha(\mu, \mathcal{V})$, assuming that $\mathcal{V}$ is the collection of all closed balls and μ is defined as in (10). In a practical realization, we start with a finite collection $\mathcal{V}_N \subset \mathcal{V}$, of N balls B_i, $i = 1, \dots, N$ which contain at least $\lfloor n(1-\alpha)+1 \rfloor$ points. In is natural to assume that if we want more than N balls, then we just add new ones to the collection $\mathcal{V}_N$, i.e,

$$\mathcal{V}_{N_1} \subset \mathcal{V}_{N_2} \qquad \text{for } N_1 < N_2. \tag{12}$$

Now for fixed N and $\mathcal{V}_N$, we intersect balls $B_i \in \mathcal{V}_N$ one by one, so that the k-the step we have the approximate level set

$$\hat{S}_{\alpha,k} := \cap_{i=1}^{k} B_i, 1 \leq k \leq N. \tag{13}$$

From the assumption (12) it follows that

$$\hat{S}_{\alpha,N} \supseteq \hat{S}_{\alpha,N+1} \supseteq \cdots \supseteq S_\alpha. \tag{14}$$

In general, for a given α it can happen that $\hat{S}_{\alpha,k} = \emptyset$. Since it is computationally hard problem to determine whether or not the set (13) is empty, and also for the purpose of visualization, we need to have some points inside the balls to decide if they belong to $\hat{S}_{\alpha,k}$ or not. Let C_M be a discrete set of points in $\mathbf{R}^d$, $|C_M| = M \geq N$, such that (i) each ball B_i, $i = 1, \dots N$ of (13) contains at least one $c \in C$ and (ii) each point $c \in C$ belongs to at least one ball B_i. It is natural to assume that sets C_M are increasing with M, that is,

$$M_1 < M_2 \implies C_{M_1} \subset C_{M_2} \tag{15}$$

Let us define

$$\hat{\hat{S}}_{\alpha,k,M} := \{c \in C_M \mid c \in \hat{S}_{\alpha,k}\}, \qquad \hat{\hat{S}}_{\alpha,N,M} = \bigcap_{k=1}^{N} \hat{\hat{S}}_{\alpha,k,M}. \tag{16}$$

For a fixed N and every $M \geq N$ we have that

$$\hat{\hat{S}}_{\alpha,N,M} \subseteq \hat{\hat{S}}_{\alpha,N,M+1} \subset \hat{S}_{\alpha,N}, \tag{17}$$

and also for every $k < N$

$$\hat{\hat{S}}_{\alpha,k,M} \supseteq \hat{\hat{S}}_{\alpha,k+1,M}. \tag{18}$$

The process of finding $\hat{\hat{S}}_{\alpha,N,M}$ as defined by (16) can end at the step $K \leq N$ in two ways: (i) if $\hat{\hat{S}}_{\alpha,K,M} = \emptyset$, (ii) if $K = N$ or $\hat{\hat{S}}_{\alpha,k,M}$ remains the same non-empty set for all k such that $K \leq k \leq N$. In the case (i) we conclude that $\hat{\hat{S}}_{\alpha,N,M} = \emptyset$. In the case (ii) we have that $\hat{\hat{S}}_{\alpha,N,M} = \hat{\hat{S}}_{\alpha,K,M}$ and we accept $\hat{\hat{S}}_{\alpha,K,M}$ as an approximation to S_α defined by (11).

Remark 3.1. 1° The output of the described procedure is $\hat{\hat{S}}_\alpha$. In order to make a contour we can find a convex hull of $\hat{\hat{S}}_\alpha$ using *QuickHull* algorithm, for example. The relations (17) remain true if $\hat{\hat{S}}_\alpha$ is replaced with its convex hull.

2° Computational experiments indicate that the convex hull of $\hat{\hat{S}}_{\alpha,M,N}$ converges to exact S_α as $M, N \to \infty$. The proof of that statement would follow from (14), (16) and (17) under some additional assumptions, which will not be further elaborated in this paper.

3° The simplest way to implement the above procedure is to take $M = N = n$, and $C = \{x_1, \ldots, x_n\}$ where x_i are sample points and intersectional balls are centered at x_i. In some cases $M = N > n$ is needed. In the rest of the paper we consider only the case $M = N$. For simplicity, in the rest of the paper we will use notation S_α in the meaning of $\hat{\hat{S}}_\alpha$ unless explicitly noted otherwise. □

3.2. **Sample augmented with artificial data points.** If we have a large and dense sample of the size n, we can take $N = n$ and use balls centered in sample points. The basic Algorithm 1 of the subsection 3.3 is presented in that setup. However, setting $N = n$ may not be sufficient in estimation of level sets and especially the deepest points (Tukey's median). As an example, consider a uniform distribution in the region bounded by circles $x^2 + y^2 = r_i^2$, $r_1 = 1$ and $r_2 = 2$. It is easy to prove (see also [8]) that the depth monotonically increases from 0 outside of the larger circle, to 1/2 at the origin, which is the true and unique median. With a sample from this distribution, we will not have data points inside the inner circle, and we can not identify the median in the way proposed above.

In similar cases and whenever we have sparse data or small sample size n, we can still visually identify depth regions and center, simply by adding artificial points to the data set. Let the data set contain points $x_1, \ldots, x_n$ and let $x_{n+1}, \ldots, x_N$ be points chosen from uniform distribution in some convex domain that contains the whole data set. Then we use a modification of the described procedure in such a way that we use augmented data set (all N points) as a criterium for stopping (in cases (i) and (ii) above), but n in formulas (10) and (11) is the cardinality of original data set.

Figure 1 shows the output of ABCDepth algorithm in the example described above. By adding artificial data points, we are able to obtain an approximate position of the Tukey's median.

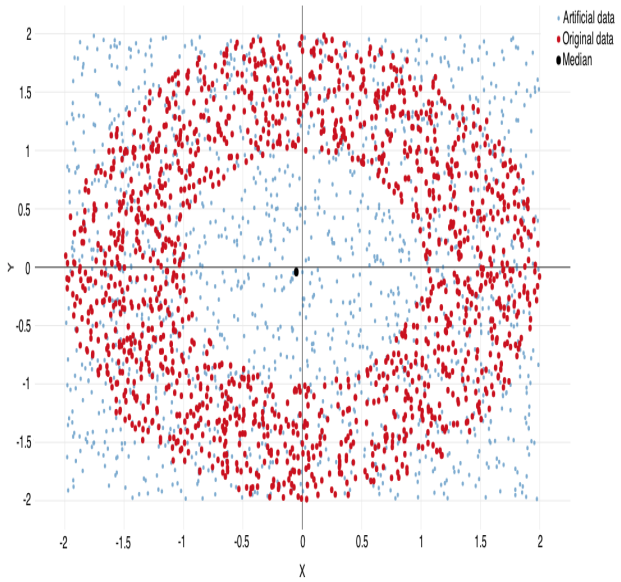

FIGURE 1. A sample from uniform distribution in a ring (red): Tukey's median (black) found with the aid of artificial points (blue).

Let us consider a triangle as in Example 2.1-4° of Section 2 with vertices $A(0,1)$, $B(-1,0)$ and $C(1,0)$. Assuming that A, B, C are sample points, all points in the interior and on the border of ABC triangle have depth $1/3$, so the depth reaches its maximum value at $1/3$. Since the original data set contains only 3 points, by adding artificial data and applying ABCDepth algorithm we can visualize the Tukey's median set as shown in Figure 2.

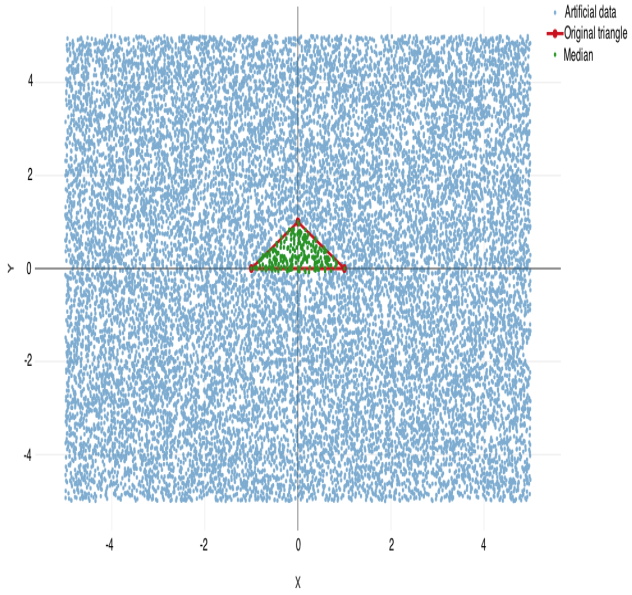

FIGURE 2. Tukey's median set of red triangle represented as triangle itself and green points inside of the triangle. Blue points are artificial data.

In the rest of this section, we describe the details of implementation of the approximate algorithm for finding Tukey's median, as well as versions of the same algorithm for finding Tukey depth of a sample point, the depth of out-of-sample point, and for data contours.

3.3. **Implementation: finding deepest points (Tukey's median).** In order to execute the calculation in (11), the first step is to construct balls for the intersection. Each ball is defined by its center and contains $\lfloor n(1-\alpha)+1 \rfloor$ nearest points, so first we calculate Euclidian inter-distances. This part of the implementation is described in lines $1-6$ of Algorithm 1. Distances are stored as a triangular matrix in a *list of lists* structure, where i-th list $(i = 1, ..., n-1)$ contains distances $d_{i+1,j}, j = 1, \ldots, i$. After sorting distances for each point, the structure that contains all n balls is populated (lines 7-10, Algorithm 1). The structure is represented as a *hashmap*, where the *key* is a center of a ball, and *value* is a *list* with $\lfloor n(1-\alpha)+1 \rfloor$ nearest points. Now, we intersect balls iteratively by increasing α by $\frac{1}{n}$.

Since this algorithm is meant to find the deepest location, there is no need to start with the minimal value of $\alpha = \frac{1}{n}$; due to Theorem 2.1, we set the initial value of α to be $\frac{1}{d+1}$. Balls intersections are shown on Algorithm 1, lines 11-17.

If the input set is sparse ABCDepth optionally creates an augmented data set of total size N as explained on page 5 and demonstrated on figures 1 and 2). Let $R_1 = \{x_1, \ldots, x_n\}$ be the original data set and let $R_2 = \{x_{n+1}, \ldots, x_N\}$ be the set of "artificial points". The algorithm creates balls with centers

in $R_1 \cup R_2$ that contain $\lfloor n(1-\alpha)+1 \rfloor$ points from R_1. The rest of the algorithm takes three phases we described above.

Data: Original data, $X_n = (\boldsymbol{x_1}, \boldsymbol{x_1}, ..., \boldsymbol{x_n}) \in \mathbb{R}^{d \times n}$
Result: List of level sets, $S = \{S_{\alpha_1}, S_{\alpha_2}, ..., S_{\alpha_m}\}$, where S_{α_m} represents a Tukey median
/* Note: S_α here means $\hat{S}_\alpha$ */

1 **for** $i \leftarrow 2$ ***to*** n **do**
2 **for** $j \leftarrow 1$ ***to*** $i-1$ **do**
3 Calculate Euclidian distance between point $\boldsymbol{x_i}$ and point $\boldsymbol{x_j}$;
4 Add distance to the list of lists ;
5 **end**
6 **end**
7 **for** $i \leftarrow 1$ ***to*** n **do**
8 Sort distances for point $\boldsymbol{x_i}$;
9 Populate structure with balls ;
10 **end**

/* Iteration Phase */
11 $size = n$, $\alpha_1 = \frac{1}{d+1}$, $k = 1$;
12 **while** $size > 1$ **do**
13 $S_{\alpha_k} = \{\bigcap_j^n B_j, |B_j| = \lfloor n(1-\alpha_k)+1 \rfloor\}$;
14 $size = |S_{\alpha_k}|$;
15 $\alpha_{k+1} = \alpha_k + \frac{1}{n}$;
16 Add S_{α_k} to S ;
17 $k = k+1$;
18 **end**

Algorithm 1: Calculating Tukey median.

The initial version of ABCDepth algorithm was presented in [2].

3.3.1. *Complexity.*

Theorem 3.1. *ABCDepth algorithm for finding approximate Tukey median has order of* $O((d+k)n^2 + n^2 \log n)$ *time complexity, where* k *is the number of iterations in the iteration phase.*

Proof. To prove this theorem we use the pseudocode of Algorithm 1. Lines 1-6 calculate Euclidian inter-distances of points. The first *for* loop (line 1) takes all n points, so its complexity is $O(n)$. Since there is no need calculate $d(x_i, x_i)$ or to calculate $d(x_j, x_i)$ if it is already calculated, the second for loop (line 2) runs in $O(\frac{n-1}{2})$ time. Finally, calculation of Euclidian distance takes $O(d)$ time. The overall complexity for lines 1-6 is:

$$O(\frac{nd(n-1)}{2}) \sim O(dn^2) \tag{19}$$

Iterating through the *list of lists* obtained in lines 1-6, the first *for* loop (line 7) runs in $O(n)$ time. For sorting the distances per each point, we use *quicksort* algorithm that takes $O(n \log n)$ comparisons to sort n points [12]. Structure populating takes $O(1)$ time. Hence, this part of the algorithm has complexity of:

$$O(n^2 \log n). \tag{20}$$

In the last phase (lines 11-18), algorithm calculates level sets by intersecting balls constructed in the previous steps. In every iteration (line 12), all n balls that contain $\lfloor n(1-\alpha_k)+1 \rfloor$ are intersected (line 13). The parameter k can be considered as a number of iterations, i.e. it counts how many times the algorithm enters in *while* loop. Each intersection has the complexity of $O(\lfloor n(1-\alpha_k)+1 \rfloor) \sim O(n)$, due to the property of the hash-based data structure we use (see for example [6]). We can conclude that the iteration phase has complexity of:

$$O(kn^2). \tag{21}$$

From (19), (20) and (21),

$$O(dn^2) + O(n^2 \log n) + O(kn^2) \sim O((d+k)n^2 + n^2 \log n), \tag{22}$$

which ends the proof.

□

Remark 3.2. 1° From the relations between S_α, $\hat{S}_\alpha$ and $\hat{\hat{S}}_\alpha$ (in notations as in 3.1, page 4), it follows that the maximal approximative depth of for a given point can not be greater its than its exact depth.

2° Under the assumption that data points are in the general position, the exact sample maximal depth is $\alpha_m = \frac{m}{n}$, where m is not greater than $\lceil \frac{n}{2} \rceil$ (see [7, Proposition 2.3]), and so by remark 1° the number k of steps satisfies the inequality

$$\frac{k-1}{n} \leq \frac{n+1}{2n} - \frac{1}{d+1}, \tag{23}$$

and the asymptotical upper bound for k is $\frac{n}{2}$.

Remark 3.3. In the case when we add artificial data points to the original data set, n in (22) should be replaced with N, where N is the cardinality of the augmented data set. The upper bound for k in (22) remains the same.

The rates of complexity with respect to n and d of Theorem 3.1 are confirmed by simulation results presented in figures 3 and 4. Measurements are taken on simulated samples of size n from d-dimensional with expectations zero and uncorrelated marginals, with $d \in \{2, ..., 10\}$ and

$$n \in \{40, 80, 160, 320, 640, 1280, 2560, 3000, 3500, 4000, 4500, 5000, 5500, 6000, 6500, 7000\}$$

The results are averaged on 10 repetitions for each fixed pair (d, n).

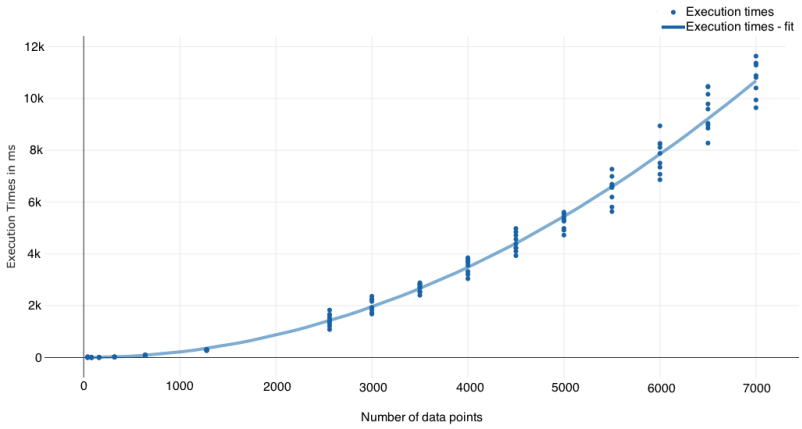

FIGURE 3. When number of points increases the execution time grows with the order of $n^2 \log n$.

3.3.2. *Examples.* Our first example is really simple and it considers n points in dimension 1 generated from normal $\mathcal{N}(0, 1)$ distribution. By running ABCDepth in this case with $n = 1000$, we get two points (as expected) in the median level set, $S_{\alpha_0.5} = \{-0.00314, 0.00034\}$. With another sample with $n = 1001$ (odd number) from the same distribution, the median set is a singleton, $S_{\alpha_0.5} = \{0.0043\}$

Now, we demonstrate data sets generated from bivariate and multivariate normal distribution.

Figure 5 and Figure 6 show the median calculated from 1000 points in dimension 2 and 3, respectively from normal $\mathcal{N}(0, 1)$ distribution. Starting from $\alpha = \frac{1}{d+1}$ the algorithm produces ~ 200 levels sets for $d = 2$ and ~ 300 level sets for $d = 3$, so not all of them are plotted. On both figures the median is represented as a black point with depth $\frac{499}{1000}$ on Figure 5, i.e. $\frac{493}{1000}$ on Figure 6.

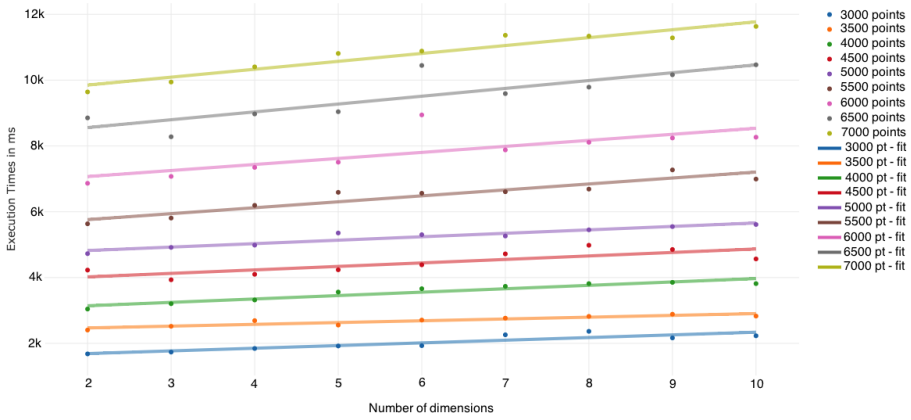

FIGURE 4. The execution time grows linearly with dimensionality.

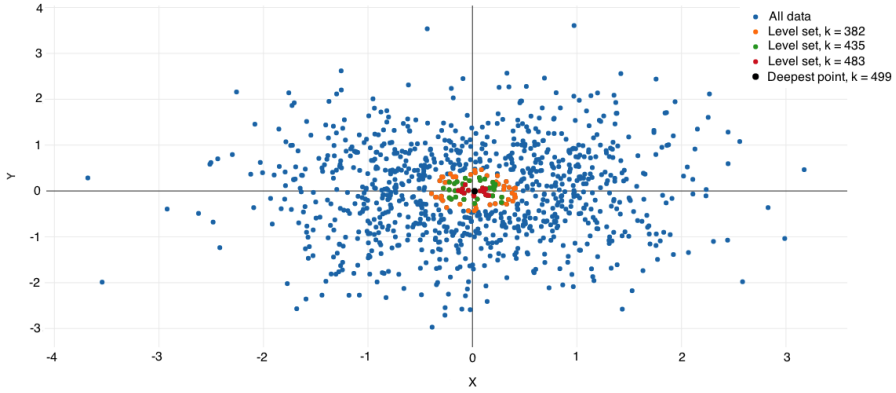

FIGURE 5. Bivariate normal distribution - four level sets, where the black point at the center is the deepest point.

All data generators that we use in this paper in order to verify and plot the algorithm output were presented at [1] and they are available within an open source project at `https://bitbucket.org/antomripmuk/generators`.

As a real data example, we take a data set which is rather sparse. The data set is taken from [27], and it has been used in several other papers as a benchmark. It contains 23 four-dimensional observations in period from 1966 to 1967 that represent seasonally adjusted changes in auto thefts in New York city. For the sake of clarity, we take only two dimensions: percent changes in manpower, and seasonally adjusted changes in auto thefts. The data is downloaded from `http://lib.stat.cmu.edu/DASL/Datafiles/nycrimedat.html`. Figure 7 shows the output of ABCDepth algorithm if we consider only points from the sample (orange point). Obviously, the approximate median belongs to the original data set. Then, we run ABCDepth algorithm with 1000 artificial data points from the uniform distribution as explained in Section 2 and earlier in this section. The approximate median obtained by this run (green point) has the same depth of $\frac{9}{23}$ as the median calculated using DEEPLOC algorithm [27] by running their Fortran code (red point). We check depths of those two points (green and red) applying *depth* function based on [24] and implemented in R "depth" package [10]. Evidently, the median, in this case, is not a singleton,

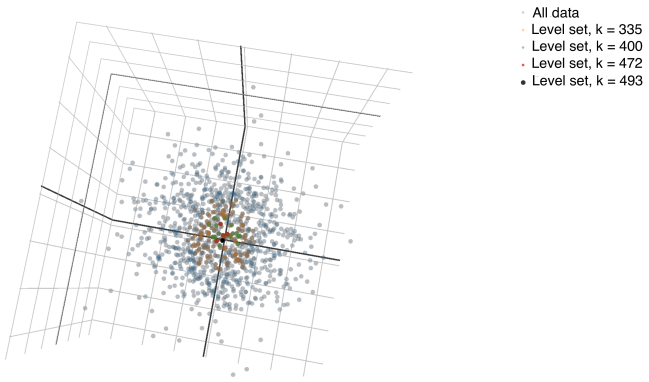

FIGURE 6. 3D normal distribution - four level sets, where the black point at the center is the deepest point.

i.e. there is more than one point with depth $\frac{9}{23}$. By adding more than 1000 artificial points, we can get more than one median point. We will discuss this example again in Section 3.5.

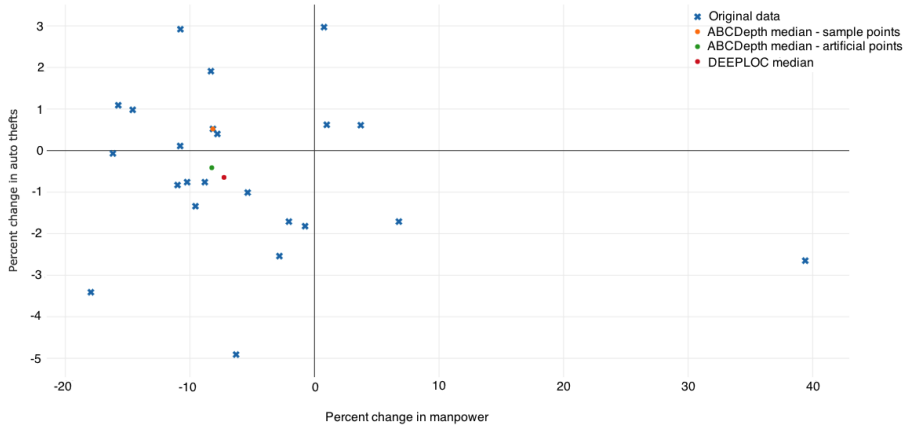

FIGURE 7. NY crime data set, comparison of Tukey medians using ABCDepth and DEEPLOC.

Another two examples are chosen from [25]. Figure 8 shows 27 two-dimensional observations that represent animals brain weight (in g) and the body weight (in kg) taken from [19]. In order to represent the same data values, we plotted the logarithms of those measurements as they did in [25].

Figure 9 considers the weight and the cost of 23 single-engine aircraft built between $1947 - 1979$. This data set is taken from [11].

As in Figure 7, in those two figures the orange point is the median obtained by running ABCDepth algorithm using only sample data. Green and red points represent outputs of ABCDepth algorithm applied by adding 1000 artificial data points from the uniform distribution and DEEPLOC median, respectively. These two examples show the importance of out-of-sample points in finding the depth levels and Tukey's median.

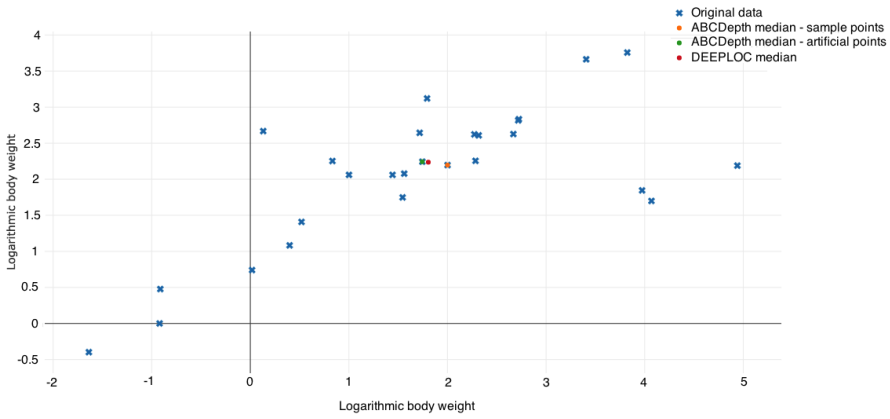

FIGURE 8. Animals data set, comparison of Tukey medians using ABCDepth and DEEPLOC.

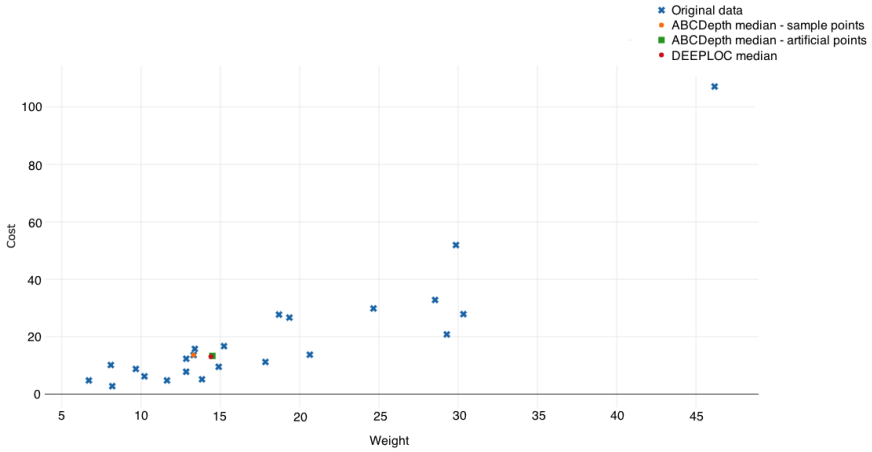

FIGURE 9. Aircraft data set, comparison of Tukey medians using ABCDepth and DEEPLOC.

3.4. **Adapted Implementation: finding the Tukey's depth of a sample point and out-of-sample point.** Let us recall that by Corollary 2.1, a point $\boldsymbol{x}$ has depth h if and only if $\boldsymbol{x} \in S_\alpha$ for $\alpha \le h$ and $\boldsymbol{x} \notin S_\alpha$ for $\alpha > h$. With a sample of size n, we can consider only $\alpha = \frac{k}{n}$, $k = 1, \dots, n$, because for $\frac{k-1}{n} < \alpha < \frac{k}{n}$, we have that $D(\boldsymbol{x}) \ge \alpha \iff D(\boldsymbol{x}) \ge \frac{k}{n}$. Therefore, the statement of Corollary 2.1 adapted to the sample distribution can be formulated as (using the fact that $S_\beta \subset S_\alpha$ for $\alpha < \beta$):

$$D(\boldsymbol{x}) = \frac{k}{n} \iff \boldsymbol{x} \in S_{\frac{k}{n}} \quad \text{and} \quad \boldsymbol{x} \notin S_{\frac{k+1}{n}}. \tag{24}$$

From (24) we derive the algorithm for Tukey's depth of a sample point $\mathbf{x}$ as follows. Let $\alpha_k = \frac{k}{n}$. The level set S_{α_1} contains all points in the sample. Then we construct S_{α_2} as an intersection of n balls that contain $n-1$ sample points. If $\boldsymbol{x} \notin S_{\alpha_2}$, we conclude that $D(\boldsymbol{x}) = 1/n$, and stop. Otherwise, we iterate this procedure till we get the situation as in right side of (24), when we conclude that the depth is $\frac{k}{n}$. The output of the algorithm is k.

Remark 3.4. As in Remark 3.2, it can be shown that the approximate depth k/n is never greater than the true depth.

Implementation-wise, in order to improve the algorithm complexity, we do not need to construct the level sets. It is enough to count balls that contain point $\boldsymbol{x}$. The algorithm stops when for some k there

exists at least one ball (among the candidates for the intersection) that does not contain $\boldsymbol{x}$. Thus, the depth of the point $\boldsymbol{x}$ is $k-1$.

With a very small modification, the same algorithm can be applied to a point $\boldsymbol{x}$ out of the sample. We can just treat $\boldsymbol{x}$ as an artificial point, in the same way as in previous sections. That is, the size of the required balls has to be $n-k+1$ points from the sample, not counting $\boldsymbol{x}$. The rest of the algorithm is the same as in the case of a sample point $\boldsymbol{x}$.

In both versions (sample or out-of-sample) we can use additional artificial points to increase the precision. The sample version of the algorithm is detailed below.

Data: Original data, $X_n = (\boldsymbol{x_1}, \boldsymbol{x_1}, ..., \boldsymbol{x_n}) \in \mathbb{R}^{d\times n}$, $\boldsymbol{x} = \boldsymbol{x}_i$ for a fixed i - the data point whose depth is calculated.

Result: Tukey depth at $\boldsymbol{x}$.

```
   /* Iteration Phase                                                    */
1  S_{α_1} = {x_1, ..., x_n};
2  for k ← 2 to n do
3  |  p = 0 - Number of balls that contain x. Its initial value is 0 ;
   |  /* Find balls that contain point x                                */
4  |  for i ← 1 to n do
5  |  |  if x ∈ B_i, where B_i contains n − k + 1 original data points then
6  |  |  |  p = p + 1;
7  |  |  end
8  |  end
9  |  if p ≠ n then
10 |  |  return k − 1
11 |  end
12 end
```

Algorithm 2: Calculating Tukey depth of a sample point.

3.4.1. *Complexity.*

Theorem 3.2. *Adapted ABCDepth algorithm for finding approximate Tukey depth of a sample point has order of $O(dn^2 + n^2 \log n)$ time complexity.*

Proof. Balls construction for Algorithm 2 is the same as in Algorithm 1 (lines 1-10), so by Theorem 3.1 this part runs in $O(dn^2 + n^2 \log n)$ time. For the point with the depth α_k algorithm enters in iteration loop k times and it iterates through all n points to find the balls that contain $\boldsymbol{x}$ point, so the whole iteration phase runs in $O(kn)$ time.

Overall complexity of the Algorithm 2 is:

$$O(dn^2 + n^2 \log n) + O(kn) \sim O(dn^2 + n^2 \log n). \tag{25}$$

□

Remark 3.5. *When the input data set is sparse or when the sample set is small, we add artificial data to the original data set in order to improve the algorithm accuracy. In that case, n in (25) should be replaced with N.*

3.4.2. *Examples.* To illustrate the output for the Algorithm 2, we use the same real data sets as we used for Algorithm 1. For all data sets we applied Algorithm 2 in two runs; first time with sample points only and second time with additional 1000 artificial points generated from uniform distribution. Points depths are verified using *depth* function from [24] implemented in [10]. For each data set we calculate the accuracy as $\frac{100k}{n}\%$, where n is the sample size and k is the number of points that has the correct depth compared with algorithm presented in [24].

On Figure 10 we showed NY crime points depths with accuracy of 26%, but if we add more points to the original data set as we showed on Figure 11, the accuracy is greatly improved to 87%.

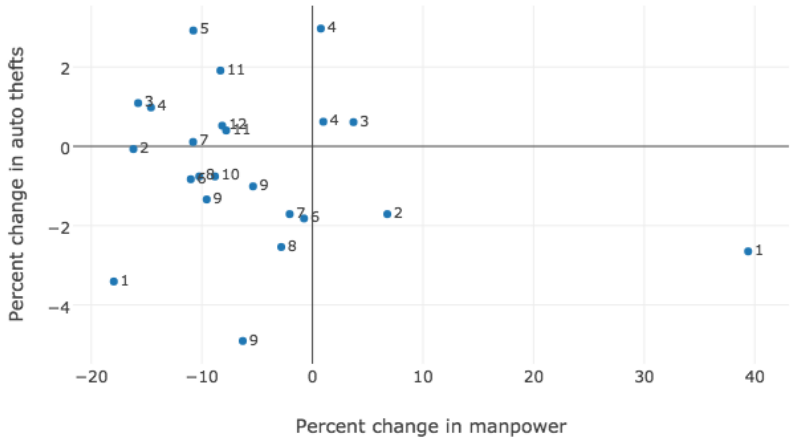

FIGURE 10. NY crime data - point depths using only original data.

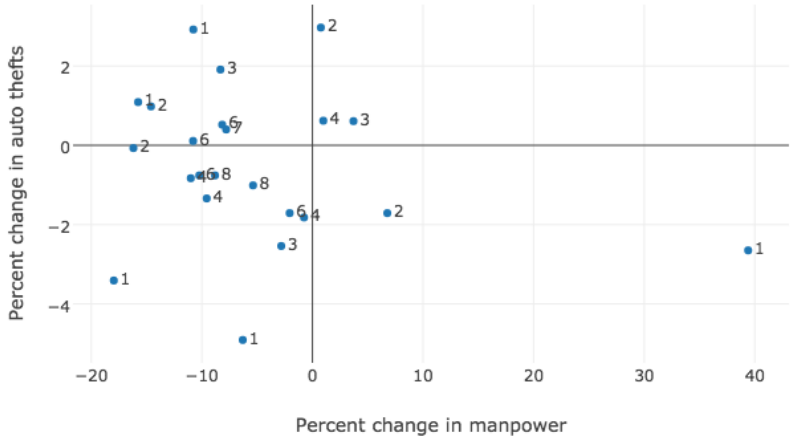

FIGURE 11. NY crime data - point depths using original and artificial data.

Figure 12 shows the same accuracy of 26% for animals data set, in the case when Algorithm 2 is run with sample points only. By adding more points as in Figure 13, the accuracy is improved to 92%.

The third example is aircraft data set presented on Figure 14 and Figure 15. The accuracy with artificial points is 95%, otherwise it is 18%.

As the last example of this section, we would like to calculate depths of the points plotted on Figure 7 using ABCDepth Algorithm 2. On Figure 7 we plotted Tukey median for NY crime data set using Algorithm 1 with artificial data points (green point) and compared the result with the median obtained by DEEPLOC (red point). Both points are out of the sample. On Figure 16 we show depths of all sample points including the depths of two median points all attained by ABCDepth Algorithm 2. Algorithm presented in [24] and ABCDepth Algorithm 2 calculate the same depth value for both median points.

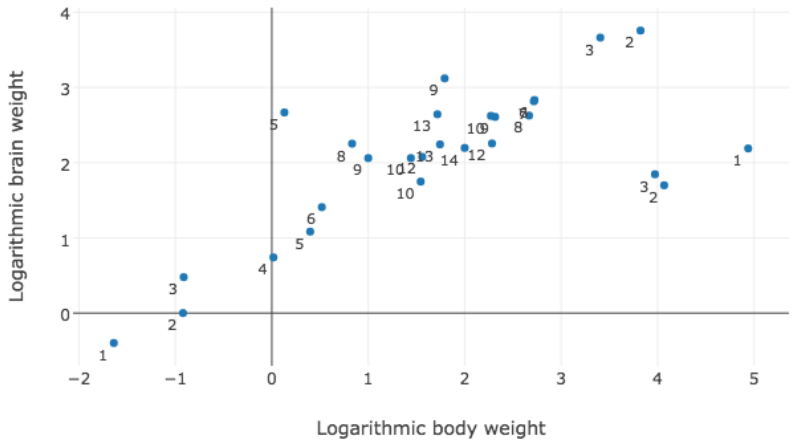

FIGURE 12. Animals data - point depths using only original data.

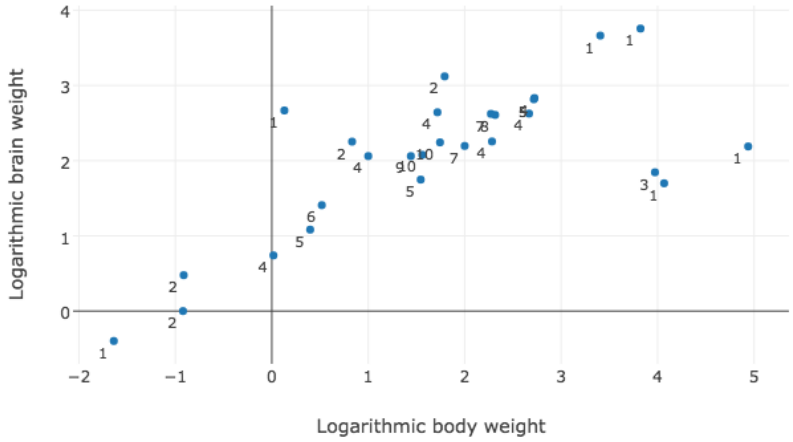

FIGURE 13. Animals data - point depths using original and artificial data.

3.5. **Adapted Implementation: finding depth contours.** Based on level sets $S = \{S_{\alpha_1}, S_{\alpha_2}, ..., S_{\alpha_m}\}$, one can obtain data depth contours applying *QuickHull* algorithm implemented in [3] on each level set.

In this purpose, to the original data set, $X_n = \{\boldsymbol{x_1}, \boldsymbol{x_2}, ..., \boldsymbol{x_n}\}$, we add artificial data points generated from the uniform distribution, $X' = \{\boldsymbol{x_{n_1}}, ..., \boldsymbol{x_N}\}$, so we denote the input data set as $X_n = \{\boldsymbol{x_1}, \boldsymbol{x_2}, ..., \boldsymbol{x_N}\}$.

As in Algorithm 2, level set S_{α_k} and consequently its depth contour D_{α_k} contain points with depth $\frac{k}{n}$, where $k = 1, ..., m$ and S_{α_m} is the deepest level set, i.e. D_{α_m} is corresponding deepest contour. Thus, in kth iteration each ball contains $n - k + 1$ points from the original data set, although the algorithm constructs and intersects all N balls. This algorithm differs from Algorithm 1 and Algorithm 2 only in iteration phase. In addition, in lines 15-18 of the Algorithm 3, convex hulls for each level set can be calculated using *QuickHull* algorithm [3].

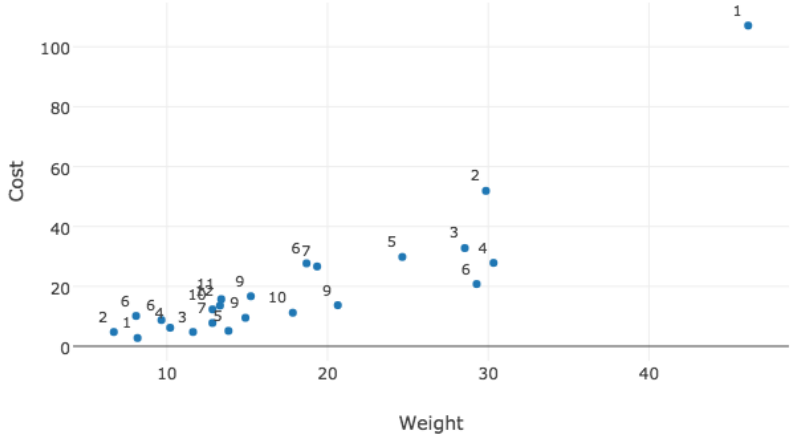

FIGURE 14. Aircraft data - point depths using only original data.

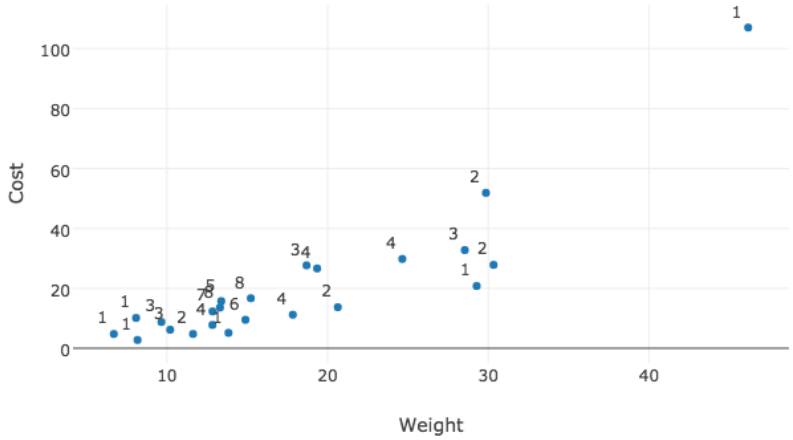

FIGURE 15. Aircraft data - point depths using original and artificial data.

Remark 3.6. From Remark 3.2 it follows that the true number of contours is never smaller than the one produced by our algorithm.

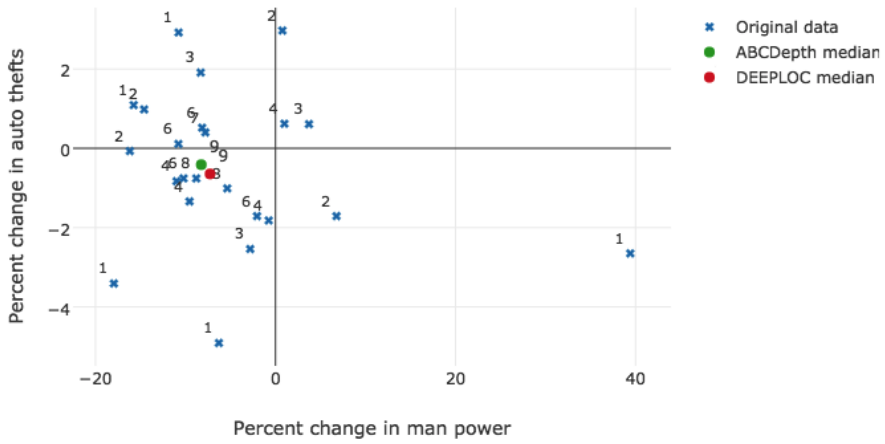

FIGURE 16. Aircraft data - point depths using original and artificial data.

Data: $X_N = (\boldsymbol{x_1}, \boldsymbol{x_2}, ..., \boldsymbol{x_n}, \boldsymbol{x_{n+1}}, ..., \boldsymbol{x_N}) \in \mathbb{R}^{d \times N}$, where x_i belongs to the original data set for $1 \leq i \leq n$ and x_i belongs to the artificial data set for $n+1 \leq i \leq N$
l - criterion for adding artificial data
p - number of artificial data points to add
Result: List of level sets, $S = \{S_{\alpha_1}, S_{\alpha_2}, ..., S_{\alpha_m}\}$, set of depth contours $D = \{D_{\alpha_1}, D_{\alpha_2}, ..., D_{\alpha_m}\}$

/* Iteration Phase */
$size = N$, $k = 1$ **while** $size > 1$ **do**
  $S_{\alpha_k} = \{\bigcap_j^N B_j, |B_j|_n = n - k + 1$ - the size w.r.t. original points only $\}$;
  $size = |S_{\alpha_k}|$;
  if $size < l$ **then**
    Generate p artificial data points located in the region of S_{α_k} ;
    Add p data points to the input set X_N ;
    $N = N + p$;
    Repeat lines 1-6 from Algorithm 1 for new N;
    Repeat lines 7-10 from Algorithm 1 for new N ;
    $S_{\alpha_k} = \{\bigcap_j^N B_j, |B_j| = n - k + 1$ - the size w.r.t. original points only $\}$;
  end
  $k = k + 1$;
  Add S_{α_k} to S ;
end
for $i \leftarrow 1$ ***to*** m **do**
  Calculate convex hull, D_{α_i}, from level set S_{α_i} ;
  Add D_{α_i} to D ;
end

Algorithm 3: Calculating depth contours.

Whenever the level set S_{α_k} contains less than l data points, the algorithm adds p artificial points to the input data set at the region of S_{α_k} in such a way that S_{α_k} is located centrally with respect to the additional artificial data. After that, the algorithm repeats the whole procedure for constructing balls for the augmented data set. The user can define the values of l and p.

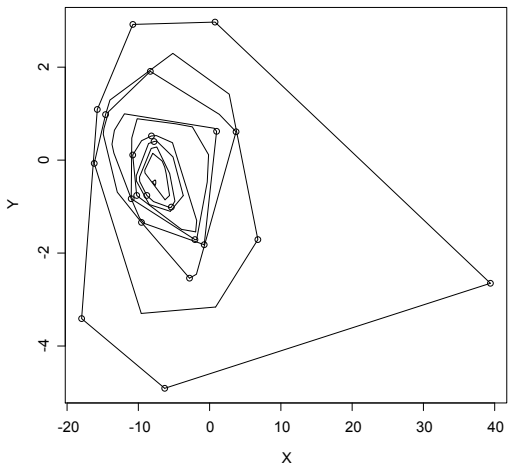

FIGURE 17. NY crime data - ISODEPTH contours.

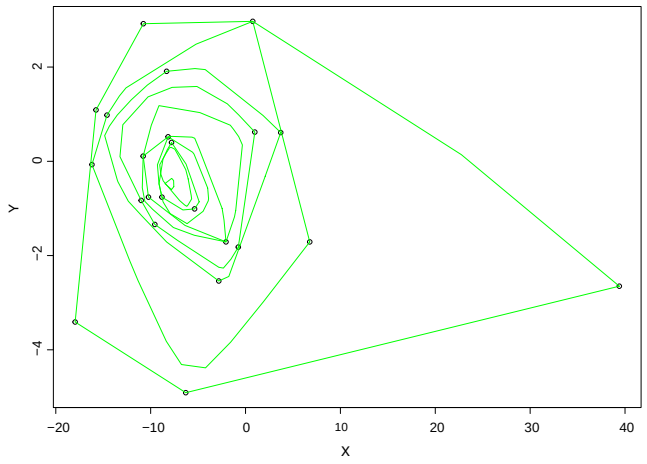

FIGURE 18. NY crime data - ABCDepth contours.

3.5.1. *Complexity discussion.* Level sets calculation has complexity which is linear in d in all three ABCDepth algorithms. That is the consequence of the fact that the number of dimensions plays a role only in Euclidian inter-distances calculation (see lines 1-6 of the Algorithm 1). Based on level sets produced in lines 1-15 of Algorithm 3, for each level set S_{α_k}, where $k = 1, ..., m$, the algorithm calculates corresponding depth contour D_{α_k} using *QuickHull* algorithm [3]. According to *QuickHull* algorithm, in $d \leq 3$ it runs in $O(n \log r)$ time, where n is the number of input set points, r is the number of processed points and it is proportional to the number of vertices in the output. Hence, the complexity ABCDepth algorithm for constructing depth contours in $d \leq 3$ is linear in d. For $d \geq 4$ the complexity of *QuickHull* grows exponentially with the number of dimensions we will not cover that case in this paper.

3.5.2. *Examples.* As a demonstration of Algorithm 3, we consider real data sets: NY crime data, data from [19] and [11]. For each data set we run *isodepth* function based on ISODEPTH algorithm [25] implemented in [10] and compare outputs.

Figure 17 and Figure 18 present contours obtained by ISODEPTH algorithm and Algorithm 3, respectively. One can note that both figures has the same number of contours, i.e. the maximal depth is $\frac{10}{23}$, although DEEPLOC yields $\frac{9}{23}$ as the approximate maximal depth. Contoure D_2 on Figure 18 contains a point $(0.76, 2.97)$ which obviously doesn't belong to the depth contour D_2 since its depth is $\frac{1}{23}$.

Contours for animals data set are shown for both algorithms on Figure 19 and Figure 20. ISODEPTH algorithm produces 11 contours, although Figure 7 in [25] has 10 contours. Algorithm 3 finds 10 contours as well, i.e. the maximal depth is $\frac{10}{27}$. Aproximate maximal depth calculated by DEEPLOC algorithm is $\frac{12}{27}$.

Figure 21 and Figure 22 present the depth contours for aircraft data set for both algorithms, ISODEPTH and Algorithm 3. Each plot contains 10 contours, i.e. the maximal depth is $\frac{10}{23}$ and DEEPLOC finds the deepest point on the same depth.

The contours produced by Algorithm 3 are similar to the contours attained by ISODEPTH algorithm for all data sets we tested, but one should keep in mind that Algorithm 3 contours are approximate and they result from intersections of balls, which in real cases can not end with straight lines with finally many balls. In most cases, a depth contour obtained by Algorithm 3 contains the exact depth contour.

4. Performance and Comparisons

According to Theorem 3.1, the complexity for calculating Tukey median grows linearly with dimension and in terms of a number of data points, it grows with the order of $n^2 \log n$. Rousseeuw and Ruts in [21] pioneered with an exact algorithm called HALFMED for Tukey median in two dimensions that runs in

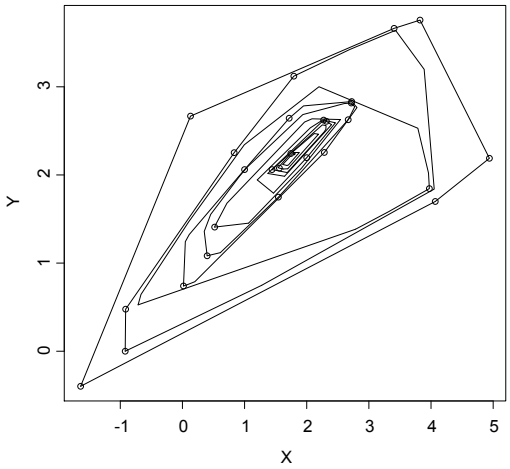

FIGURE 19. Animals data - ISODEPTH contours.

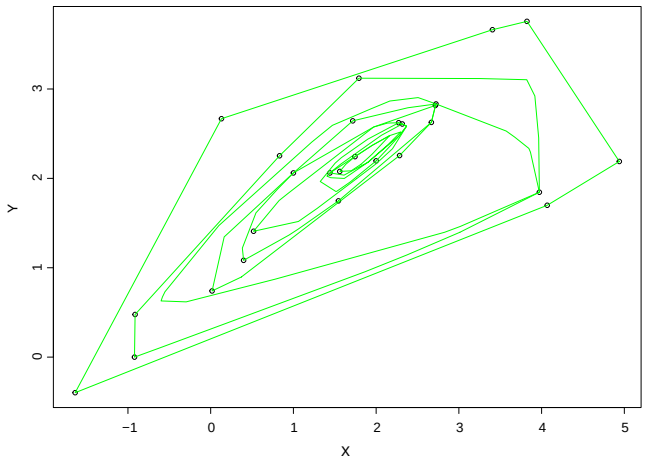

FIGURE 20. Animals data - ABCDepth contours.

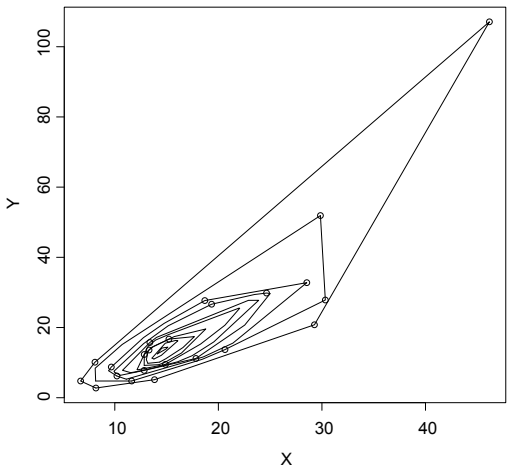

FIGURE 21. Aircraft data - ISODEPTH contours.

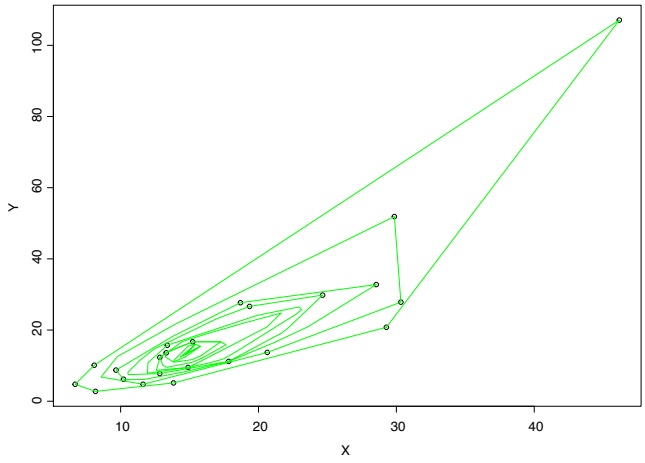

FIGURE 22. Aircraft data - ABCDepth contours.

$O(n^2 log^2 n)$ time. This algorithm is better than ABCDepth for $d = 2$, but it processes only bivariate data sets. Struyf and Rousseeuw in [27] implemented the first approximate algorithm called DEEPLOC for finding the deepest location in higher dimensions. Its complexity is $O(kmn \log(n + kdn + md^3 + mdn))$ time, where k is the number of steps taken by the program and m is the number of directions, i.e. vectors constructed by the program. This algorithm is very efficient for low-dimensional data sets, but for high-dimensional data sets ABCDepth algorithm outperforms DEEPLOC. Chan in [4] presents an approximate randomized algorithm for maximum Tukey depth. It runs in $O(n^{d-1})$ time and it is not implemented yet.

In Table 1 execution times of DEEPLOC algorithm and ABCDepth algorithm for finding Tukey median are reported. The measurements are performed using synthetic data generated from the multivariate normal $\mathcal{N}(0, 1)$ distribution. In this table, we demonstrate how ABCDepth algorithm behaves with thousands of high-dimensional data points. It takes $\sim$ 13 minutes for $n = 7000$ and $d = 2000$. Since DEEPLOC algorithm doesn't support data sets with $d > n$ and returns the error message: "the dimension

TABLE 1. Compare DEEPLOC and ABCDepth execution times in seconds.

d	Algorithm	n												
		320	640	1280	2560	3000	3500	4000	4500	5000	5500	6000	6500	7000
50	Deeploc	4.43	7.15	12.65	23.87	30.93	31.79	37.66	45.35	50.72	63.13	63.75	84.13	69.61
	ABCDepth	0.15	0.63	2.86	4.95	7.27	8.65	12.51	14.18	17.51	22.18	25.86	29.24	37.34
100	Deeploc	19.42	22.85	33.81	77.45	69.04	105.56	97.39	120.05	140.04	131.85	127.36	212.42	183.27
	ABCDepth	0.22	0.92	2.03	7.83	9.78	13.14	17.89	23.52	30.6	39.18	49.03	68.46	82.02
500	Deeploc	-	1616.53	*	*	*	*	*	*	*	*	*	*	*
	ABCDepth	0.693	3.181	8.4	27.9	41.61	53.73	71.95	89.36	109.22	140.18	151.45	180.5	213.01
1000	Deeploc	-	-	*	*	*	*	*	*	*	*	*	*	*
	ABCDepth	1.165	3.99	14.389	54.18	74.38	98.73	129.85	164.96	203.37	246.54	286.17	344.94	39.16
2000	Deeploc	-	-	-	*	*	*	*	*	*	*	*	*	*
	ABCDepth	2.21	7.86	27.25	107.46	132.77	180.02	243.1	297.6	386.75	475.87	554.23	666.4	764.74

TABLE 2. Aveerage time per data point.

d	n												
	320	640	1280	2560	3000	3500	4000	4500	5000	5500	6000	6500	7000
50	0.07	0.21	1.21	8.23	12.64	19.22	28.56	42.04	64.33	77.45	98.79	121.86	150.73
100	0.08	0.25	1.23	8.18	13.91	20.48	28.51	44.31	65.55	81.91	99.96	123.84	154.65
500	0.13	0.42	1.84	11.42	17.93	21.42	35.41	52.07	73.21	95.18	119.82	141.88	176.12
1000	0.17	0.53	2.52	13.53	20.13	32.35	41.71	58.72	82.92	103.84	138.69	155.32	200.55
2000	0.26	0.94	4.12	18.32	28.12	38.79	56.04	73.79	102.98	124.48	156.54	186.59	232.45

should be at most the number of objects", we denoted those examples with $-$ sign in the table. The sign $*$ means that the median is not computable at least once in 12 hours.

ABCDepth algorithm for finding Tukey depth of a point runs in $O(dn^2 + n^2 \log n)$ as we showed in Theorem 3.2. Most of the algorithms for finding Tukey depth are exact and at the same time computationally expensive. One of the first exact algorithms for bivariate data sets, called LDEPTH, is proposed by Rousseeuw and Ruts in [20]. It has complexity of $O(n \log n)$ and like HALFMED it outperforms ABCDepth for $d = 2$. Rousseeuw and Struyf in [24] implemented an exact algorithm for $d = 3$ that runs in $O(n^2 \log n)$ time and an approximate algorithm for $d > 3$ that runs in $O(md^3 + mdn)$ where m is the number directions, i.e. all directions perpendicular to hyperplanes through d data points. The later work of Chen et al. in [5] presented approximate algorithms based on the third approximation method of Rousseeuw and Struyf in [24] reducing the problem from d to k dimensions. The first one, for $k = 1$, runs in $O(\epsilon^{1-d} dn)$ time and the second one, for $k \geq 2$, runs in $O((\epsilon^{-1} c \log n)^d)$, where ϵ and c are empirically chosen constants. The another exact algorithm for finding Tukey depth in $\mathbf{R}^d$ is proposed by Liu and Zuo in [14], which proves to be extremely time-consuming (see Table 5.1 of Section 5.3 in [18]) and the algorithm involves heavy computations, but can serve as a benchmark. Recently, Dyckerhoff and Mozharovskyi in [9] proposed two exact algorithm for finding halfspace depth that run in $O(n^d)$ and $O(n^{d-1} \log n)$ time.

Table 2 shows execution times of ABCDepth algorithm for finding a depth of a sample point. Measurements are derived from synthetics data from the multivariate standard normal distribution. Execution time for each data set represents averaged time consumed per data point. Most of the execution time ($\sim 95\%$) is spent on balls construction (see lines 1-10 of the Algorithm 1), while finding a point depth itself (iteration phase of the Algorithm 2) is really fast since it runs in $O(kn)$ time.

In Section 3.5 we presented ABCDepth algorithm for calculating level sets and in addition it can construct depth contours using *QuickHull* algorithm. Its complexity is linear in d for $d \leq 3$. For $d = 2$ there are two exact algorithms for constructing depth contours. The first one, called ISODEPTH, is proposed by Ruts nad Rousseeuw in [25] and for $n < 1000$ the time of the proposed algorithm behaves as a multiple of $n^2 \log n$, although according to *isodepth* function implemented in R "depth" package [10] ISODEPTH takes several minutes to calculate contours from 1000 points generated from bivariate normal distribution. The second algorithm is presented by Miller et al. in [17] which computes all bivariate depth contours in $O(n^2)$ time. For the depth contours in dimensions $d > 2$ Liu et al. proposed an algorithm in [13] that runs in $O(n^p \log n)$ time.

The ABCDepth algorithm has been implemented in Java. Tests for all algorithms are run using one kernel of Intel Core i7 (2.2 GHz) processor.

Acknowledgements

We would like to express our gratitude to Anja Struyf and coauthors for sharing the code and the data that were used in their papers of immense importance in the area. Answering to Yijun Zuo's doubts about the first arXiv version of this paper and solving difficult queries that he was proposing, helped us to improve the presentation and the algorithms. The second author acknowledges the support by grants III 44006 and 174024 from Ministry of Education, Science and Technological Development of Republic of Serbia.

References

[1] M. Bogicevic and M. Merkle. Multivariate Medians and Halfspace Depth: Algorithms and Implementation. In *Proc. 1st International Conference on Electrical, Electronic and Computing Engineering (IcETRAN 2014), Vrnjačka Banja, Serbia*, volume 1, page 27, 2014.

[2] M. Bogicevic and M. Merkle. Data Centrality Computation: Implementation and Complexity Calculation. In *Proc. 2nd International Conference on Electrical, Electronic and Computing Engineering (IcETRAN 2015), Srebrno Jezero, Serbia*, volume 1, page 23, 2015.

[3] Hannu Huhdanpaa C. Bradford Barber, David P. Dobkin. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software*, 22:469–483, 1996.

[4] T. M. Chan. An Optimal Randomized Algorithm for Maximum Tukey Depth. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 430–436. ACM, New York, 2004.

[5] Dan Chen, Pat Morin, and Uli Wagner. Absolute approximation of Tukey depth: Theory and experiments. *Comput. Geom.*, 46:566–573, 2013.

[6] Bolin Ding and Arnd Christian Knig. A Fast set intersection in memory. *Proceedings of the VLDB Endowment* , 4:255–266, 2011.

[7] D. L. Donoho and M. Gasko. Breakdown properties of location estimates based on halfspace depth and projected outlyingness. *Ann. Statist.*, 20:1803–1827, 1992.

[8] S. Dutta, A. K. Ghosh, and P. Chaudhuri. Some intriguing properties of Tukey's half-space depth. *Bernoulli*, 17:1420–1434, 2011.

[9] Rainer Dyckerhoff and Pavlo Mozharovskyi. Exact computation of the halfspace depth. *Computational Statistics and Data Analysis*, 98:19–30, 2016.

[10] Maxime Genest, Jean-Claude, and Jean-Francois Plante. Package depth, 2012.

[11] J.B. Gray. Graphics for regression diagnostics,. *ASA Proc. Statistical Computing Section*, pages 102–107, 1985.

[12] C. A. R. Hoare. Algorithm 64: Quicksort. *Comm. Acm.*, 4:321, 1961.

[13] X. Liu, K. Mosler, , and P. Mozharovskyi. Fast computation of Tukey trimmed regions in dimension $p > 2$. *arXiv:1412.5122*, 2014.

[14] X. Liu and Y. Zuo. Computing halfspace depth and regression depth. *Communications in Statistics Simulation and Computation*, 43:969–985, 2014.

[15] Milan Merkle. Jensen's inequality for medians. *Stat. Prob. Letters*, 71:277–281, 2005.

[16] Milan Merkle. Jensen's inequality for multivariate medians. *J. Math. Anal. Appl.*, 370:258–269, 2010.

[17] Kim Miller, Suneeta Ramaswami, Peter Rousseeuw, J.Antoni Sellares, Diane Souvaine, Ileana Streinu, and Anja Struyf. Efficient computation of location depth contours by methods of computational geometry. *Statistics and Computing*, 13:153–162, 2003.

[18] Pavlo Mozharovskyi. *Contributions to depth-based classification and computation of the Tukey depth.* PhD thesis, Faculty of Economics and Social Sciences, University of Cologne, 2014.

[19] P. J. Rousseeuw and A. M. Leroy. Robust Regression and Outlier Detection. *Wiley*, page 57, 1997.

[20] P. J. Rousseeuw and I. Ruts. Bivariate Location Depth. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 45:516–526, 1996.

[21] P. J. Rousseeuw and I. Ruts. Constructing the Bivariate Tukey Median. *Constructing the Bivariate Tukey Median*, 8:827–839, 1998.

[22] P. J. Rousseeuw and I. Ruts. The depth function of a population distribution. *Metrika*, 49:213–244, 1999.

[23] Peter J. Rousseeuw and Mia Hubert. Statistical depth meets computational geometry: a short survey. *arXiv*, 2015.

[24] Peter J. Rousseeuw and Anja Struyf. Computing location depth and regression depth in higher dimension. *Statistics and Computing*, 8:193–203, 1998.

[25] I. Ruts and P. J. Rousseeuw. Computing depth contours of bivariate point clouds. *Computational Statistics and Data Analysis*, 23:153–168, 1996.

[26] C. G. Small. A survey of multidimensional medians. *arXiv*, 58:263–277, 1990.

[27] A. Struyf and P. J. Rousseeuw. High-dimensional computation of the deepest location. *Comp. Statist. & Data Anal.*, 34:415–426, 2000.

[28] John Tukey. Mathematics and Picturing Data. In *Proc. International Congress of Mathematicians, Vancouver 1974*, volume 2, pages 523–531, 1975.

[29] Y. Zhou and R. Serfling. Multivariate spatial U-quantiles: A Bahadur-Kiefer representation, a Theil-Sen estimator for multiple regression, and a robust dispersion estimator. *J. Statist. Plann. Inference*, 138:1660–1678, 2008.

[30] Y. Zuo and R. Serfling. General notions of statistical depth function. *Ann. Stat.*, 28:461–482, 2000.